

Fault Tolerance in Open MPI

Joshua Hursey

Indiana University
Open Systems Lab.
jjhursey@open-mpi.org
www.cs.indiana.edu/~jjhursey

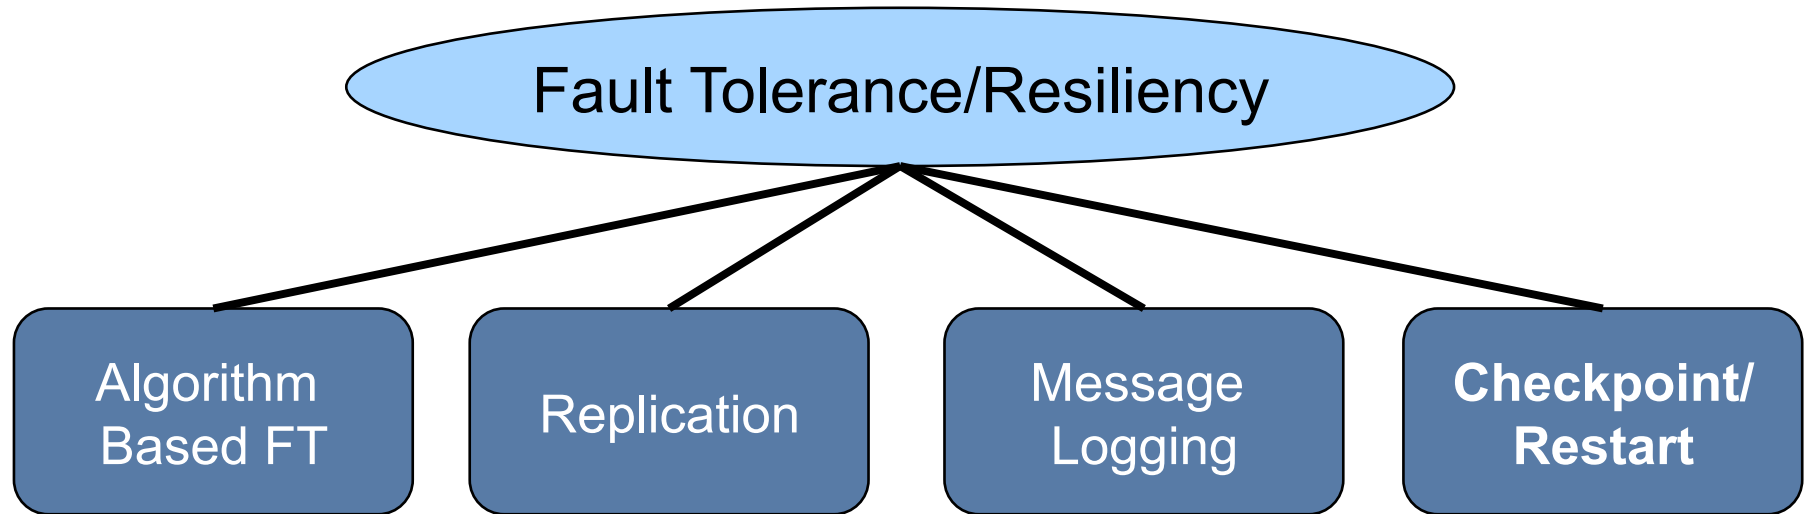
Fault Tolerance/Resiliency


Algorithm
Based FT


Replication

Message
Logging

**Checkpoint/
Restart**








High Level Goals

- Deliver usable features to end users
 - Don't publish and run
- Extensible C/R research infrastructure
 - Focused development areas
 - Apples-to-apples comparisons
 - Opportunities for public release & support



Features

- ❑ Fault Tolerance
- ❑ Debugging
- ❑ Process Migration

Infrastructure

- ❑ Checkpoint Service
- ❑ Coordination Protocol
- ❑ Runtime Coordination
- ❑ File Management
- ❑ Internal Coordination
- ❑ Recovery Service
- ❑ *In development...*

Feature: Fault Tolerance

- Transparent, checkpoint/restart driven by:
 - System Administrator
 - Resource Manager/Scheduler
 - Application

```
shell$ ompi-checkpoint 1234
```

```
Snapshot Ref.: 0 ompi_global_snapshot_1234.ckpt
```

```
shell$ ompi-checkpoint 1234
```

```
Snapshot Ref.: 1 ompi_global_snapshot_1234.ckpt
```

Sequence Numbers

Global Snapshot Reference


```
shell$ ompi-restart ompi_global_snapshot_1234.ckpt
```

Hursey, J., et. al., *The design and implementation of checkpoint/restart process fault tolerance for Open MPI*.
IEEE IPDPS, 2007.

Feature: Debugging

“My program only fails after **4 hours** when running with **>512 processes**.”

- Step-backward
(a.k.a. reverse execution)
 - Combination of checkpoint/restart and message logging
- Specified a C/R interface for:
 - Parallel debugger,
 - C/R enabled MPI implementation,
 - Checkpoint/restart service



Feature: Process Migration

Transparent process migration without residual dependencies

```
shell$ ompi-migrate --off odin001 123  
shell$ ompi-migrate --off odin001 --onto odin002,odin003 123
```

- ❑ Proactive Migration
 - Move processes when asked by predictor (e.g., CFTS FTB, RAS, ...)
- ❑ Cluster Management
 - Move processes when asked by end user
- ❑ Automatic Recovery
 - Rollback all processes to the last checkpoint, restart failed processes on new/spare resources.

Performance Impact

Latency

Interconnect	No C/R	With C/R	% Overhead
Ethernet (TCP)	49.92 μ s	50.01 μ s	0.2 %
InfiniBand	8.25 μ s	8.78 μ s	6.4 %
Myrinet MX	4.23 μ s	4.81 μ s	13.7 %
Shared Memory	1.84 μ s	2.15 μ s	16.8 %



Bandwidth

Interconnect	No C/R	With C/R	% Overhead
Ethernet (TCP)	738 Mbps	738 Mbps	0.0 %
InfiniBand	4703 Mbps	4703 Mbps	0.0 %
Myrinet MX	8000 Mbps	7985 Mbps	0.2 %
Shared Memory	5266 Mbps	5258 Mbps	0.2 %

NASA Parallel Benchmarks: 0 – 0.6 %

Gromacs (DPPC): 0%

Checkpoint Overhead





BT Class C 36 Procs
4.2 GB/120 MB

EP Class D 32 Procs
102 MB/3.2 MB

Hursey, J., et. al., *Interconnect Agnostic Checkpoint/Restart in Open MPI*. ACM HPDC, 2009.

Checkpoint Overhead





SP Class C 36 Procs
1.9 GB/54 MB

LU Class C 32 Procs
1 GB/32 MB

Hursey, J., et. al., *Interconnect Agnostic Checkpoint/Restart in Open MPI*. ACM HPDC, 2009.

Checkpoint Overhead




Gromacs (DPPC) 8 Procs
267 MB/33 MB


Gromacs (DPPC) 16 Procs
473 MB/30 MB

Hursey, J., et. al., *Interconnect Agnostic Checkpoint/Restart in Open MPI*. ACM HPDC, 2009.

Checkpoint Bottlenecks

98.8% File I/O
0.7% Modex
0.3% Coord. Protocol
0.2% Internal Coord.





Features


- ❑ Fault Tolerance
- ❑ Debugging
- ❑ Process Migration

Infrastructure

- ❑ Checkpoint Service
- ❑ Coordination Protocol
- ❑ Runtime Coordination
- ❑ File Management
- ❑ Internal Coordination
- ❑ Recovery Service
- ❑ *In development...*

Distributed Snapshots

The global state of a distributed system is defined as the *state of all processes and all connected channels* in the system.



6 processes + 9 channels

C/R Infrastructure in Open MPI

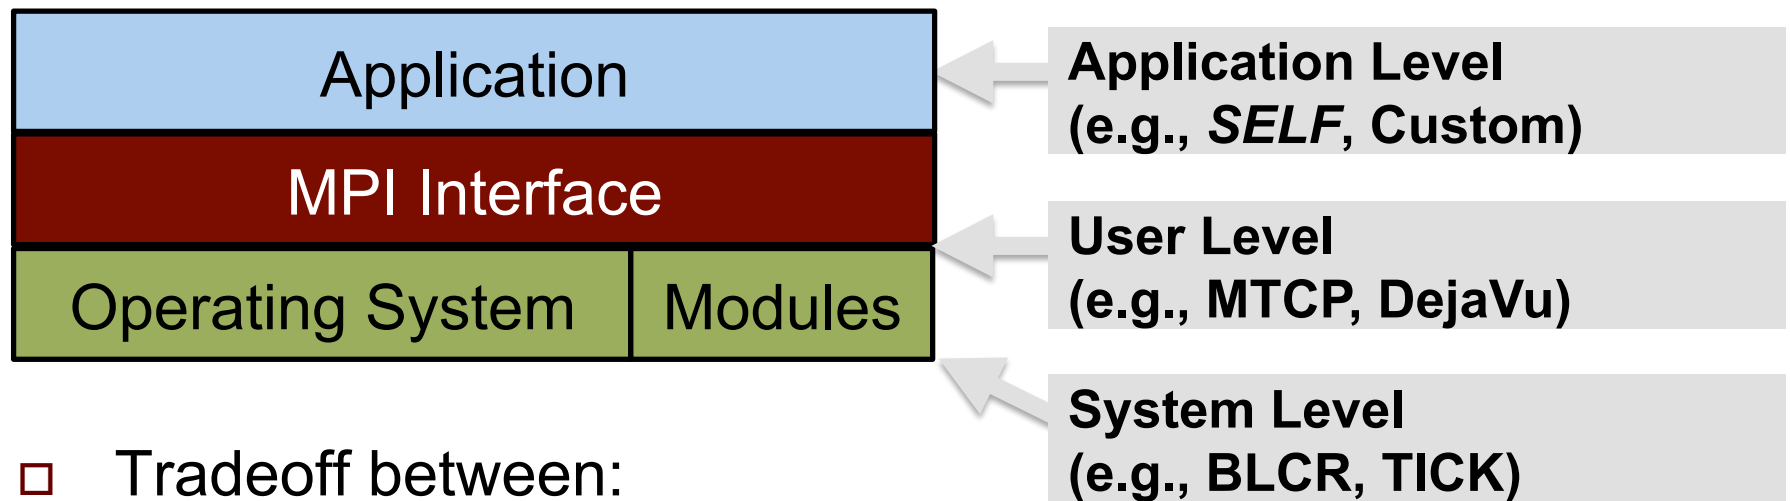
Process

CRS

Runtime

Checkpoint/Restart Service (CRS)


Capture the state of a single process



- ❑ Tradeoff between:
 - ❑ Transparency
 - ❑ Performance
 - ❑ Portability
- ❑ API and/or callbacks required for MPI support

C/R Infrastructure in Open MPI


Process



Runtime

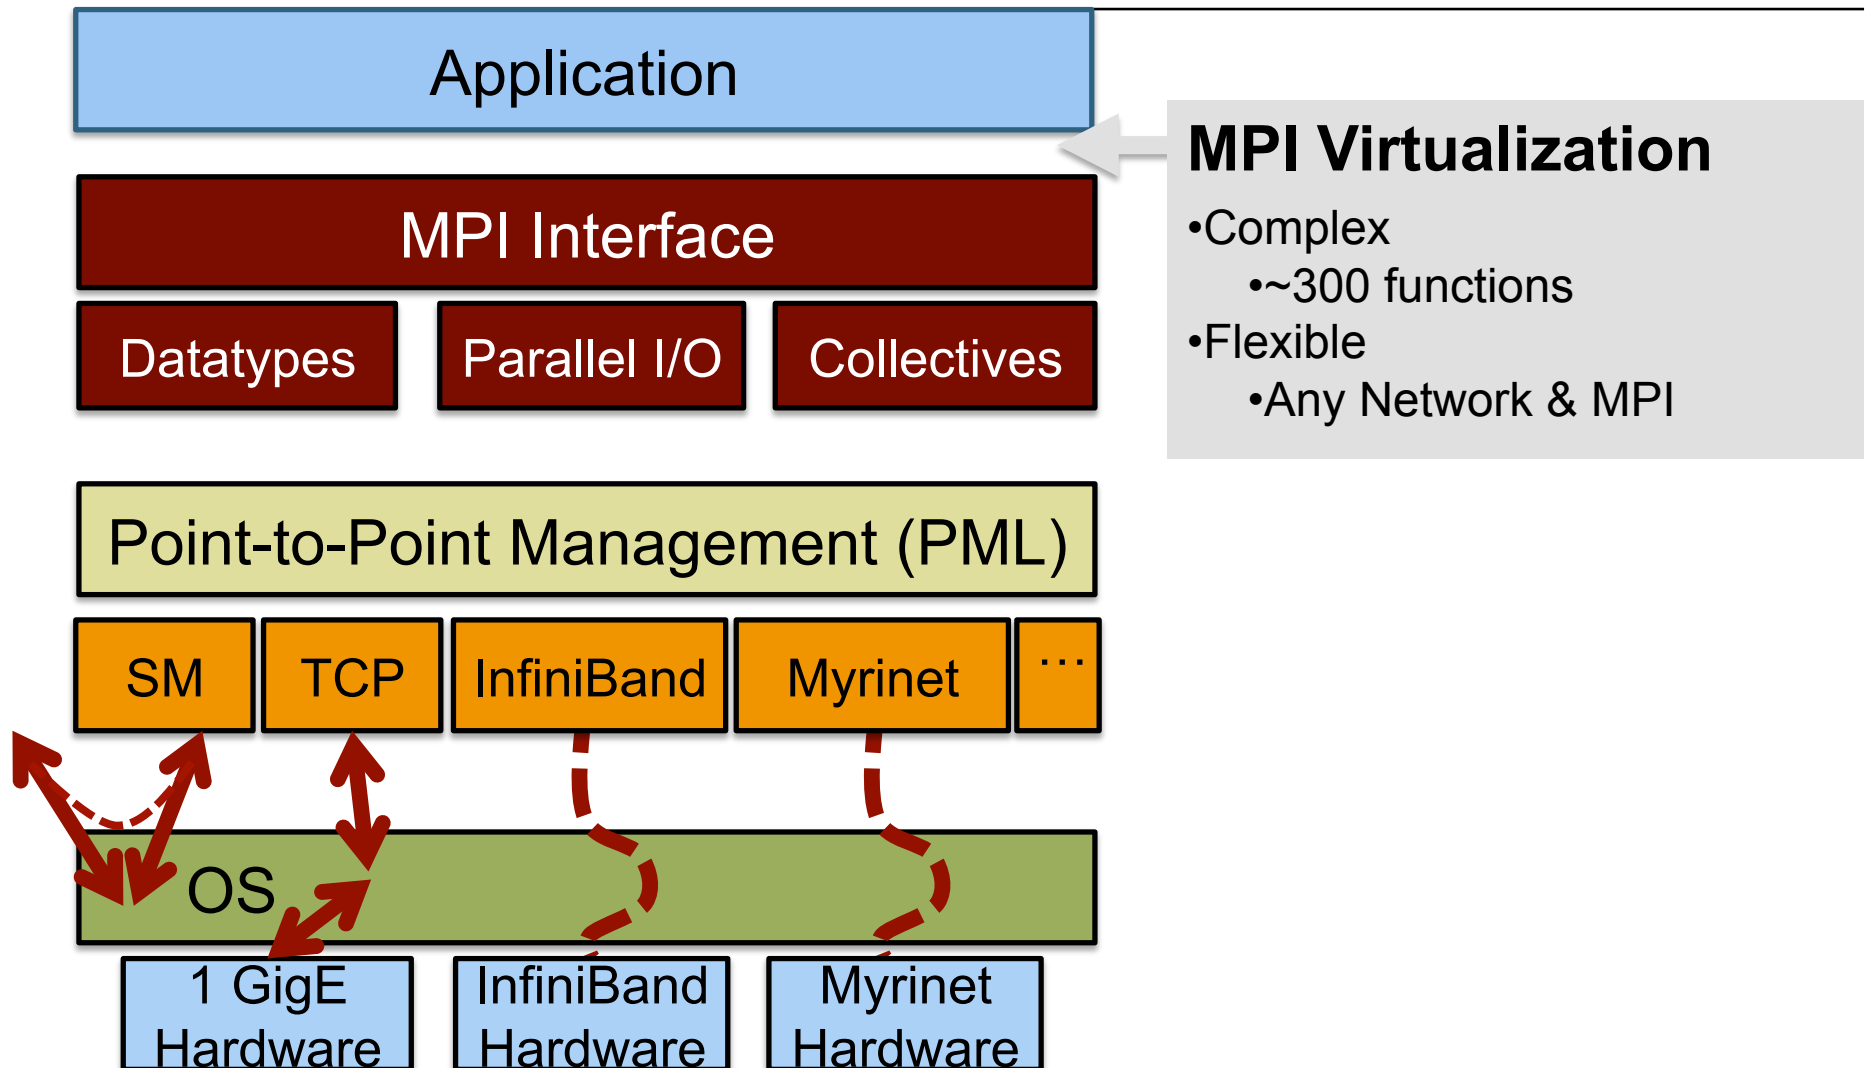
Message Coordination Protocol

Capture the state of all connected channels.
Find a (strongly) consistent state.



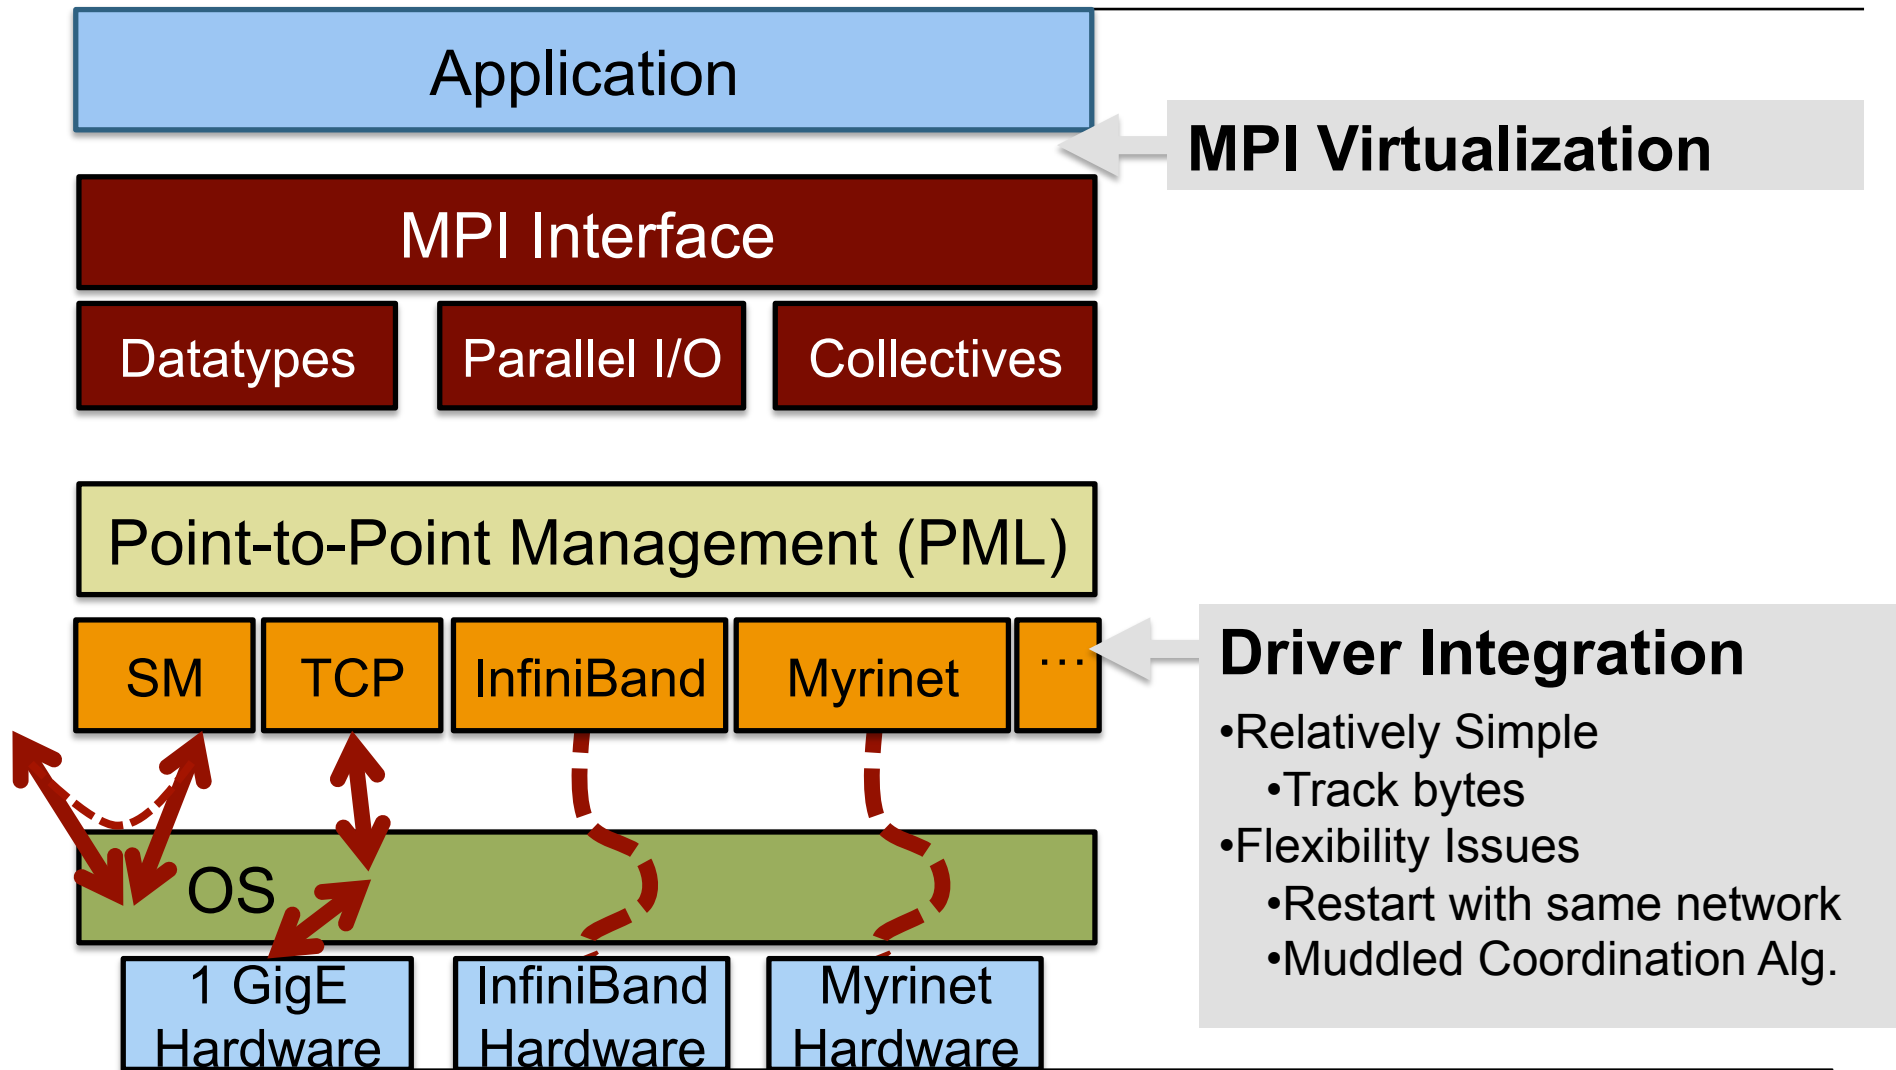
- Common Coordination Algorithms
 - Chandy/Lamport's Distributed Snapshots
 - CoCheck's Ready Message
 - LAM/MPI's Bookmark Exchange

Coordination Protocol Integration



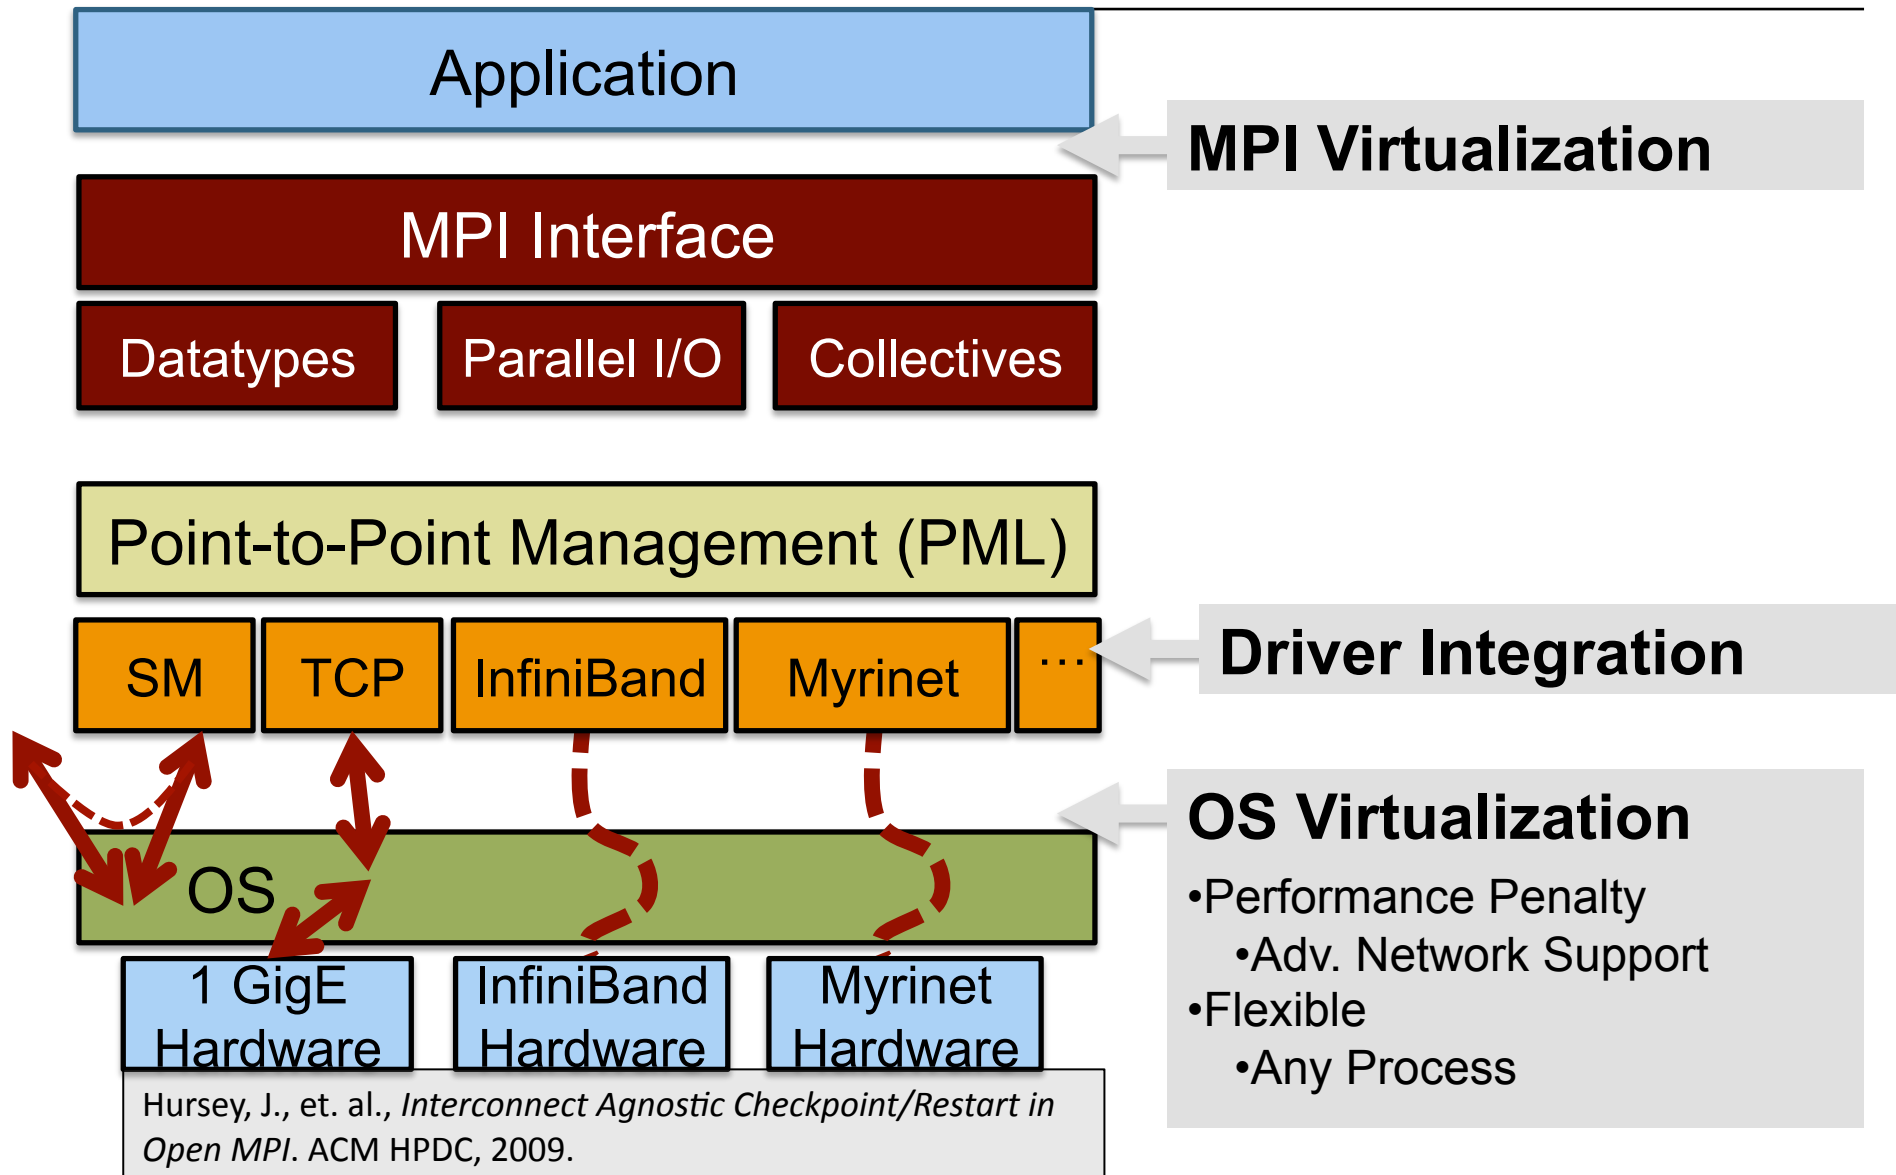
Hursey, J., et. al., *Interconnect Agnostic Checkpoint/Restart in Open MPI*. ACM HPDC, 2009.

Coordination Protocol Integration

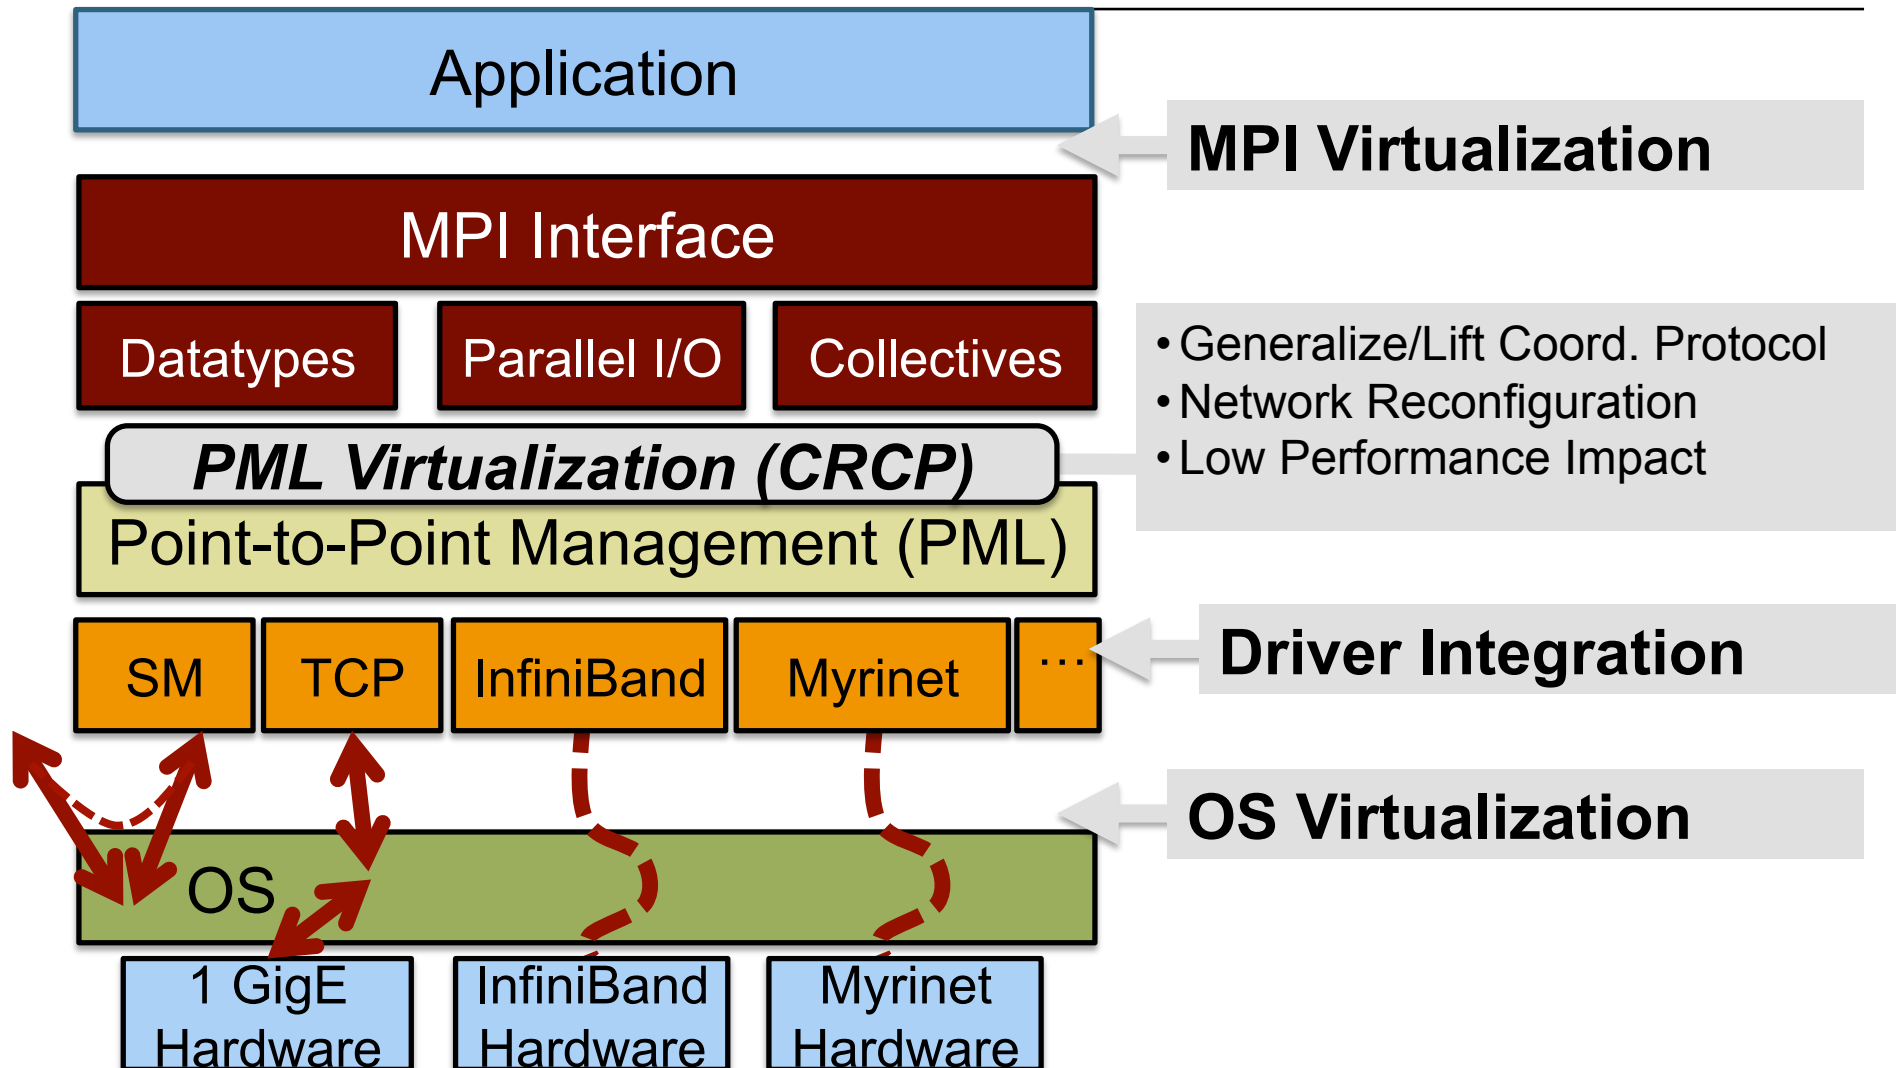


Hursey, J., et. al., *Interconnect Agnostic Checkpoint/Restart in Open MPI*. ACM HPDC, 2009.

Coordination Protocol Integration




Coordination Protocol Integration



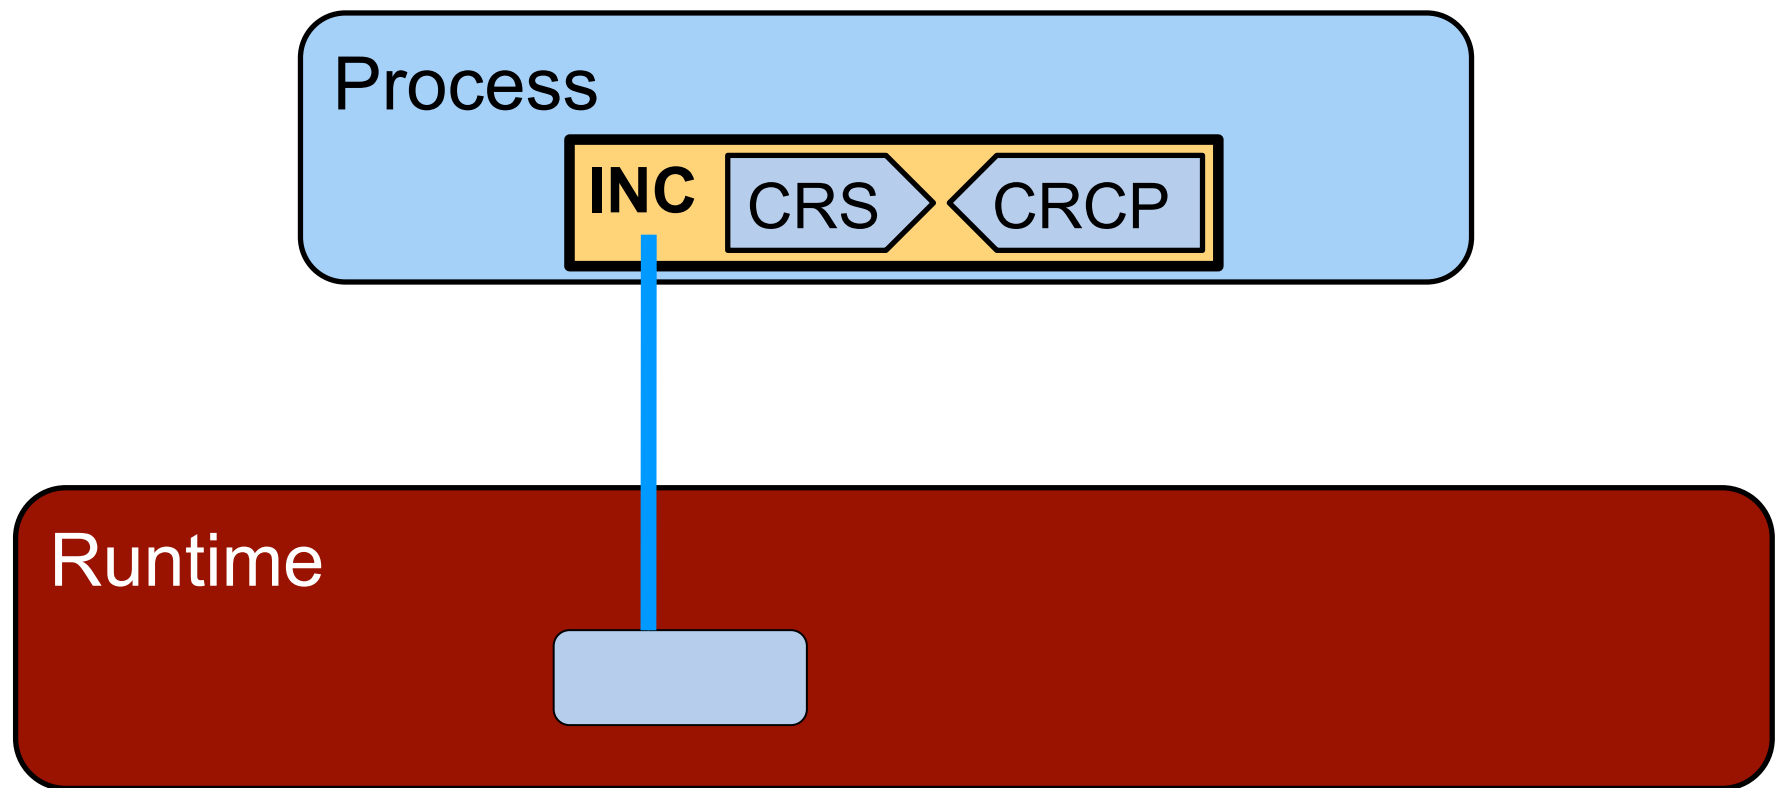
Hursey, J., et. al., *Interconnect Agnostic Checkpoint/Restart in Open MPI*. ACM HPDC, 2009.

Network Reconfiguration



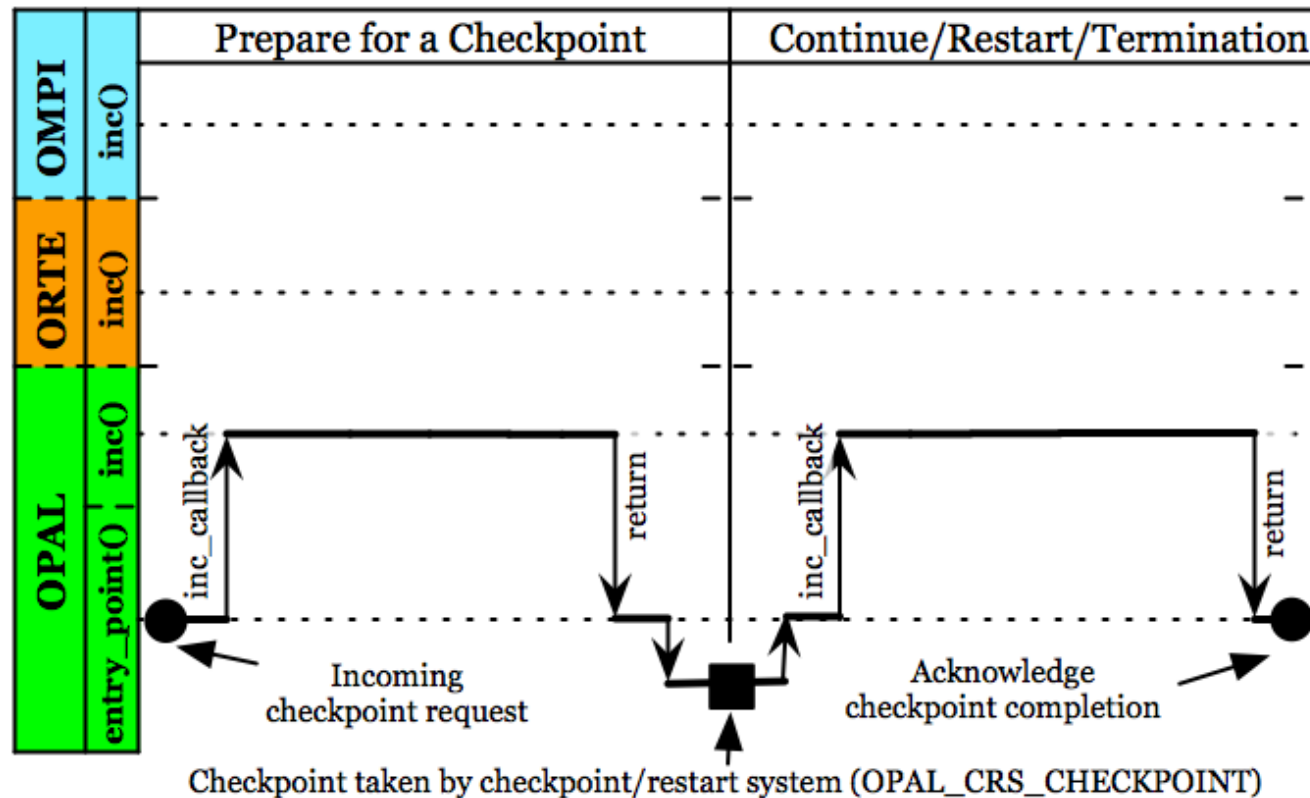
Baseline MX — Baseline OpenIB — Baseline SM — Multilevel — Baseline TCP —

C/R Infrastructure in Open MPI



Internal Coordination (INC)

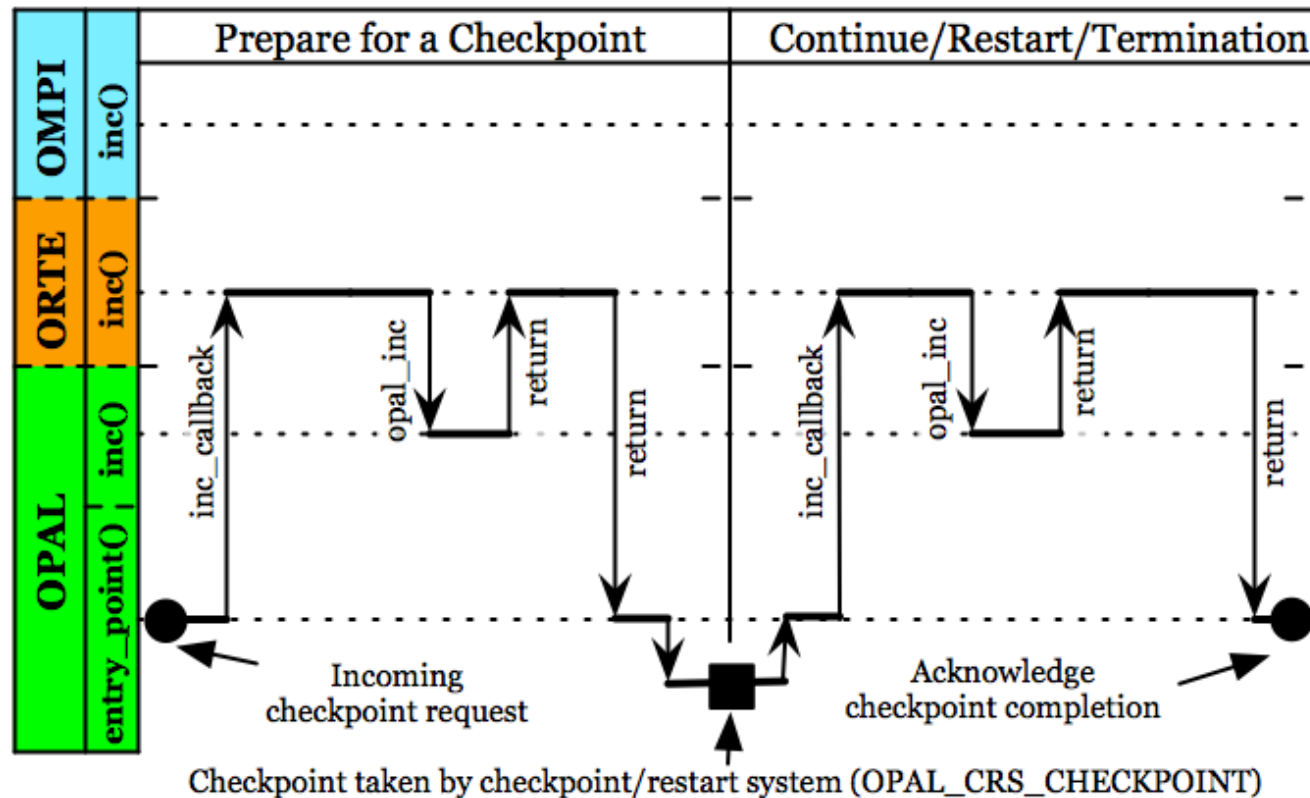
Intra-process coordination of notifications to all layers and frameworks in Open MPI



Hursey, J., et. al., *The design and implementation of checkpoint/restart process fault tolerance for Open MPI*. IEEE IPDPS, 2007.

Internal Coordination (INC)

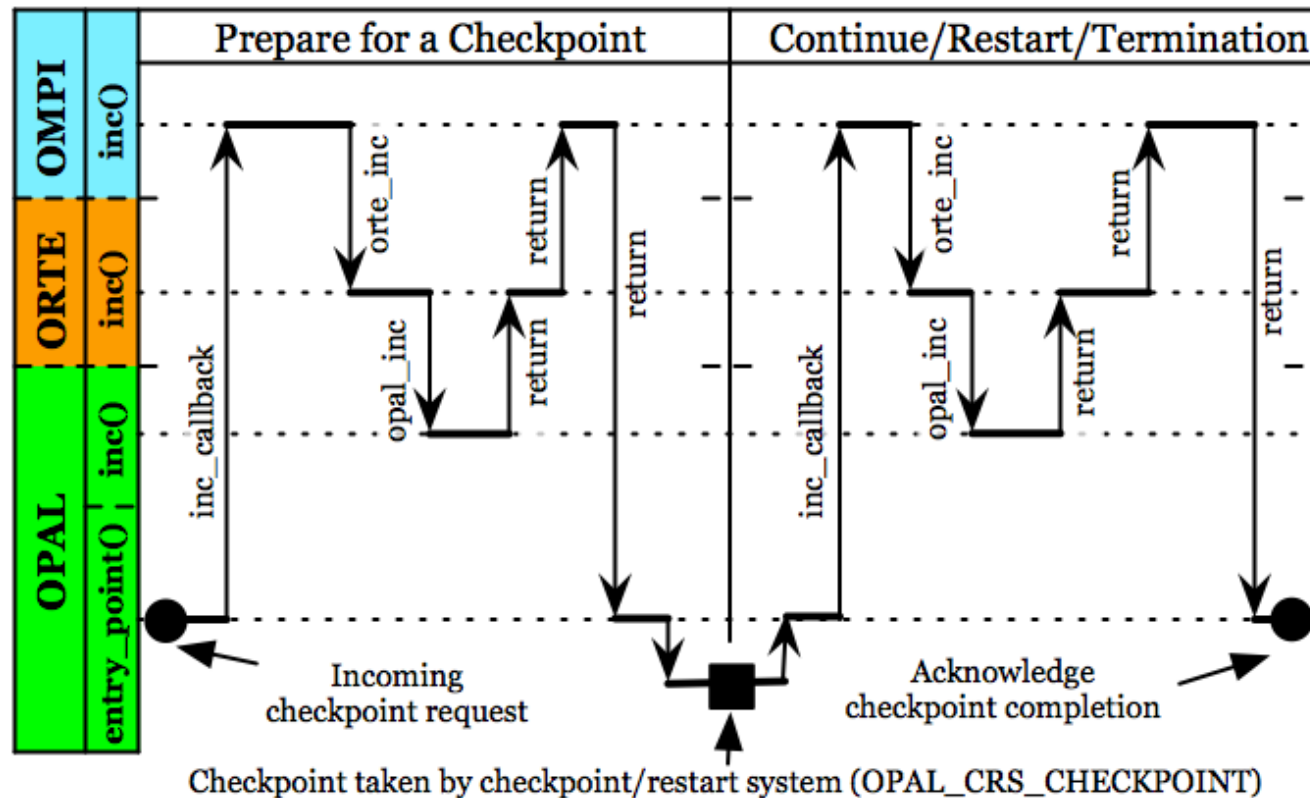
Intra-process coordination of notifications to all layers and frameworks in Open MPI



Hursey, J., et. al., *The design and implementation of checkpoint/restart process fault tolerance for Open MPI*. IEEE IPDPS, 2007.

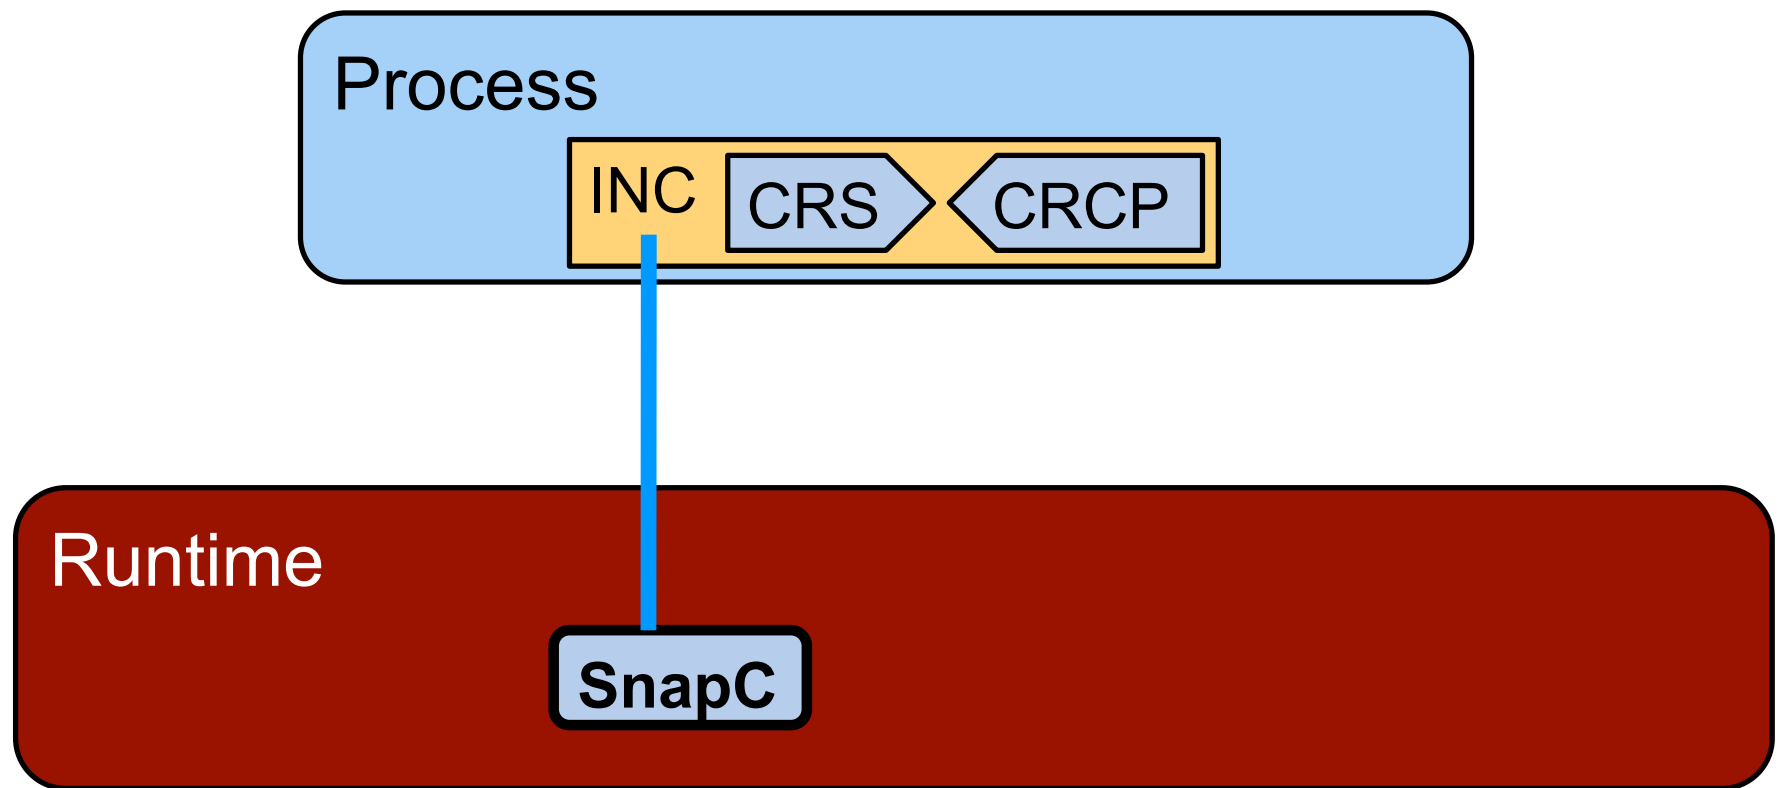
Internal Coordination (INC)

Intra-process coordination of notifications to all layers and frameworks in Open MPI



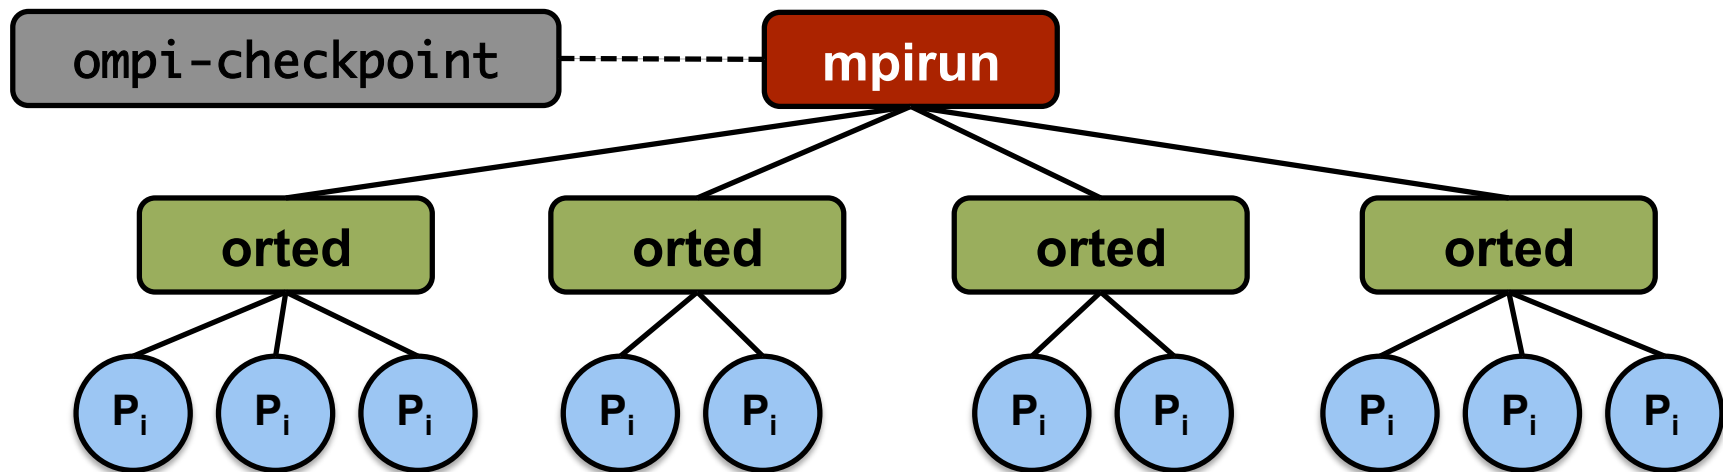
Hursey, J., et. al., *The design and implementation of checkpoint/restart process fault tolerance for Open MPI*. IEEE IPDPS, 2007.

C/R Infrastructure in Open MPI



Runtime Coordination (SnapC)

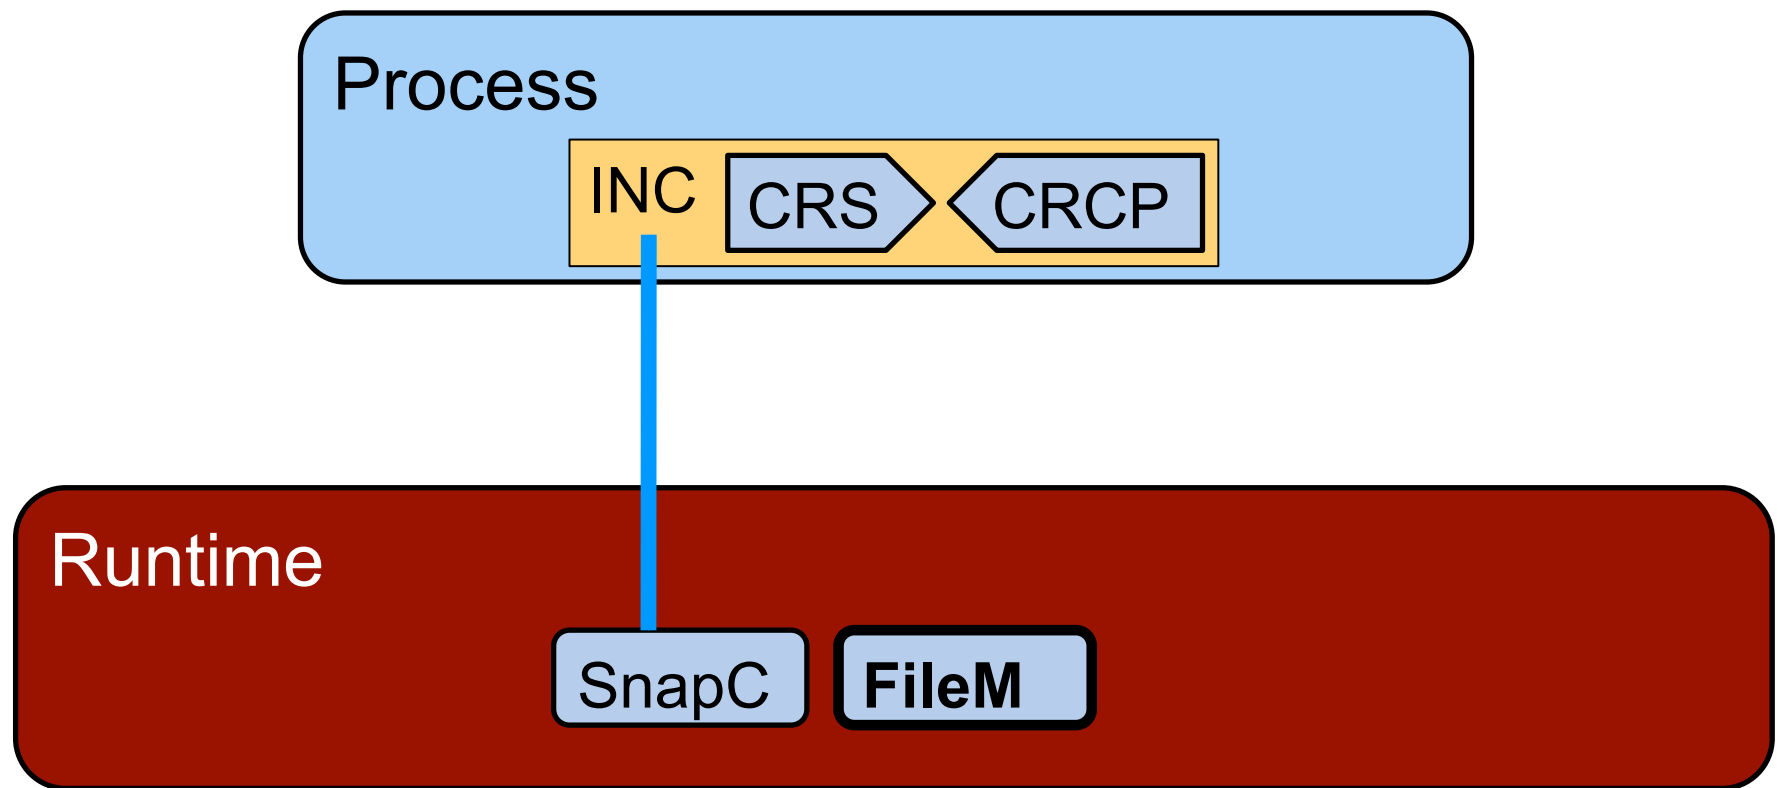
Coordinate all checkpoint related activities in ORTE,
and interact with command line tools



1. Initiate the per process local checkpoint operation
2. Monitor the progress of the checkpoint operation
3. Aggregate the local snapshots into a global snapshot
4. Preserve global snapshot on stable storage

Hursey, J., et. al., *The design and implementation of checkpoint/restart process fault tolerance for Open MPI*.
IEEE IPDPS, 2007.

C/R Infrastructure in Open MPI




File Management (FileM)

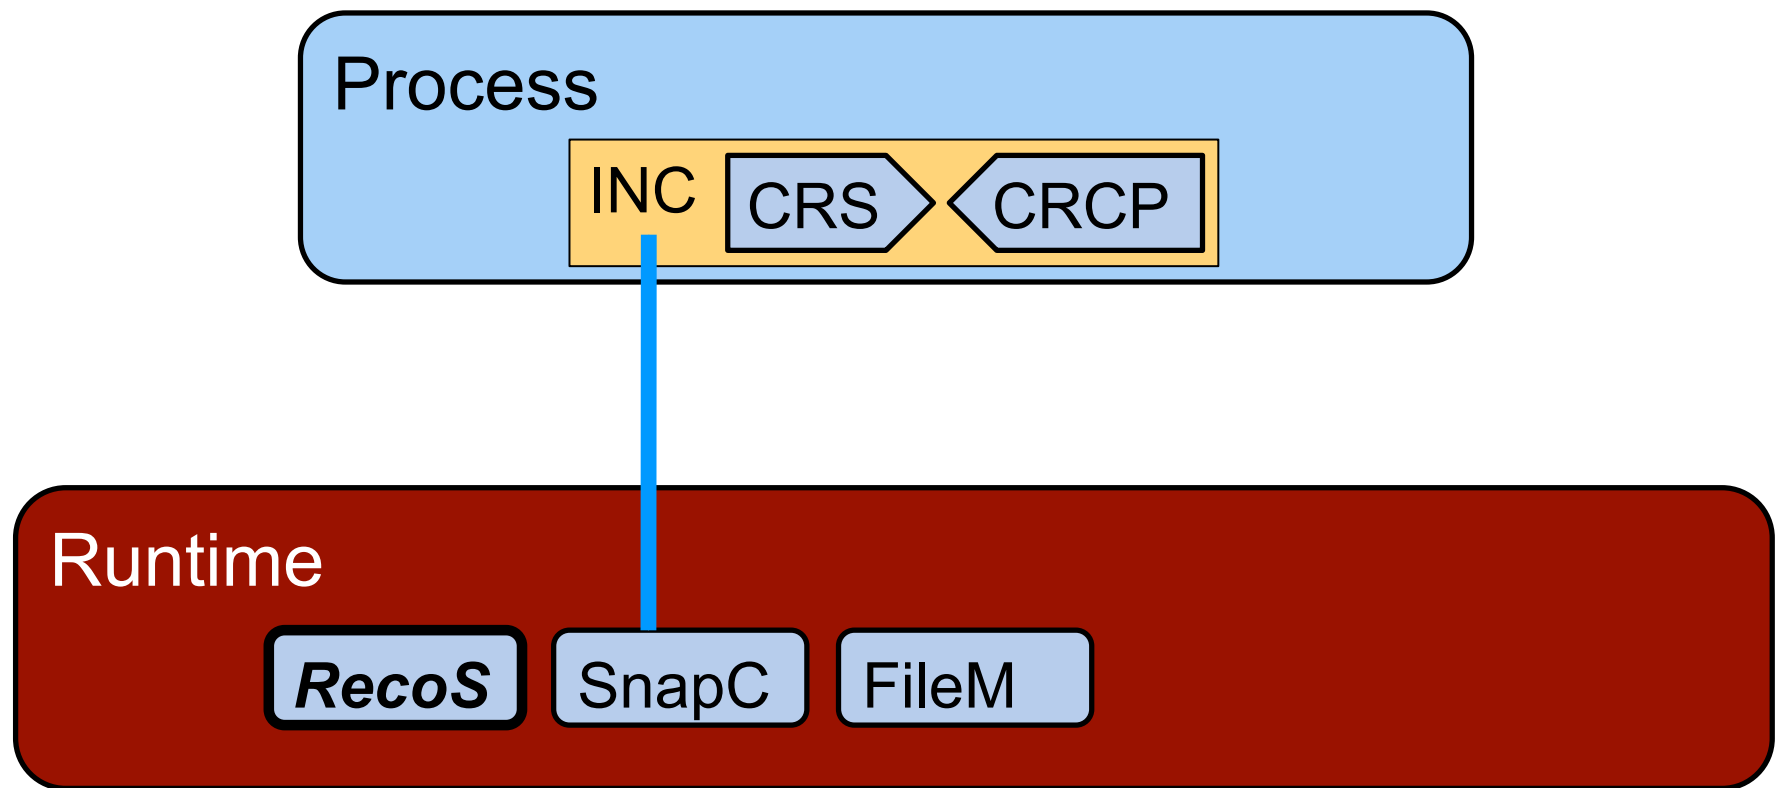
Management the movement of files from one file system to another

- **Stable Storage:**
 - *Any storage device that survives the maximum number of expected faults in a system*
- **Interface:**
 - Get() – Gather to stable storage
 - Put() – Broadcast to local storage
 - Remove() – Cleanup temporary files

File Management (FileM)



C/R Infrastructure in Open MPI

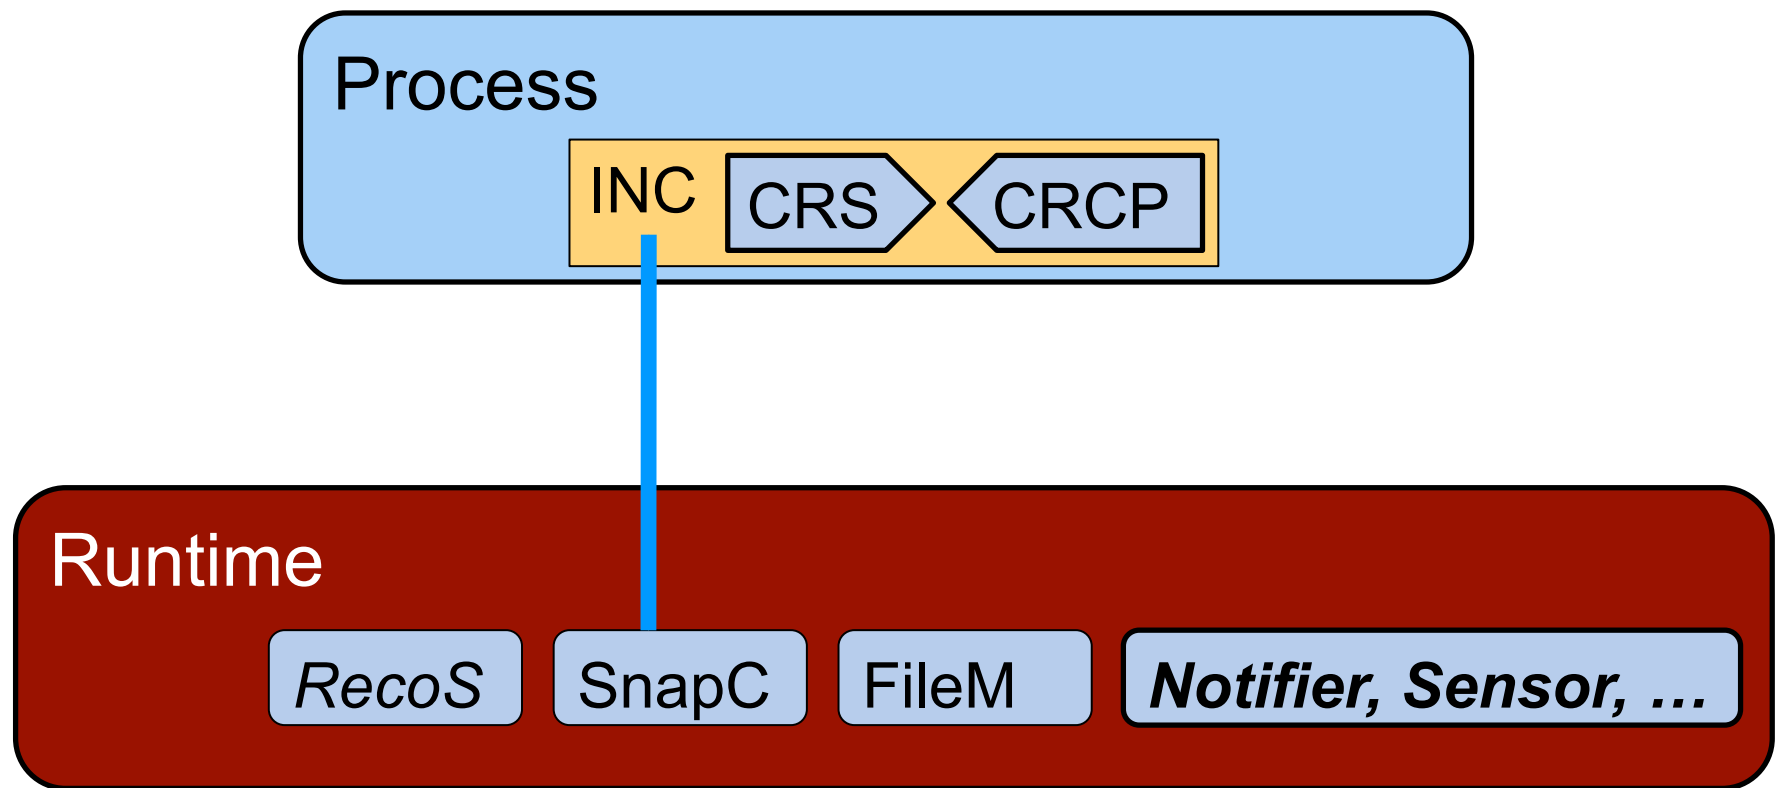


Recovery Service (RecoS)

Policy enforcement for runtime fault recovery and preventative actions


- Policy variations:
 - **Abort:**
Terminate job
 - **Ignore:**
Stabilize and run without the failed process
 - **Migrate:**
Preventatively move processes between resources
 - **Restart:**
Automatically restart from the last checkpoint
- Can be used to support MPI layer policies optionally expressed by an application

C/R Infrastructure in Open MPI



Services in Development

- Event Monitoring
 - CIFS FTB, IPMI, and others
- Fault Prediction
 - Experimenting with models using:
 - Old logs (historical perspective)
 - Current logs (recent history)
 - Hardware sensors (present environment)
- Fault Detection
 - Current is simple heartbeat mechanism
 - Looking at more scalable protocols




Features

- ❑ Fault Tolerance
- ❑ Debugging
- ❑ Process Migration

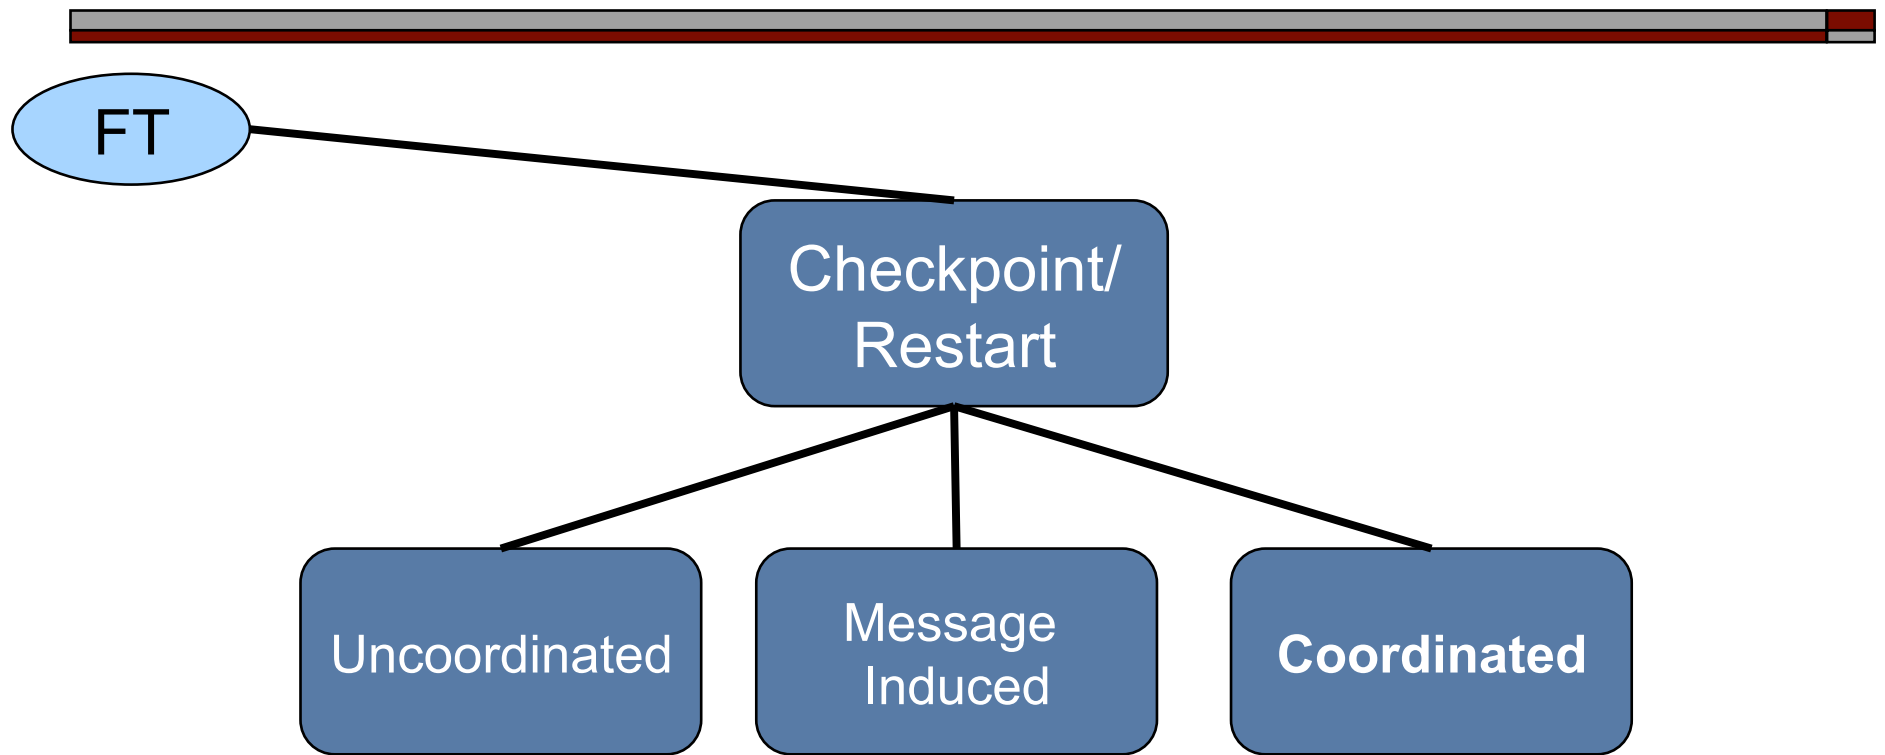
Infrastructure

- ❑ Checkpoint Service
- ❑ Coordination Protocol
- ❑ Runtime Coordination
- ❑ File Management
- ❑ Internal Coordination
- ❑ Recovery Service
- ❑ *In development...*

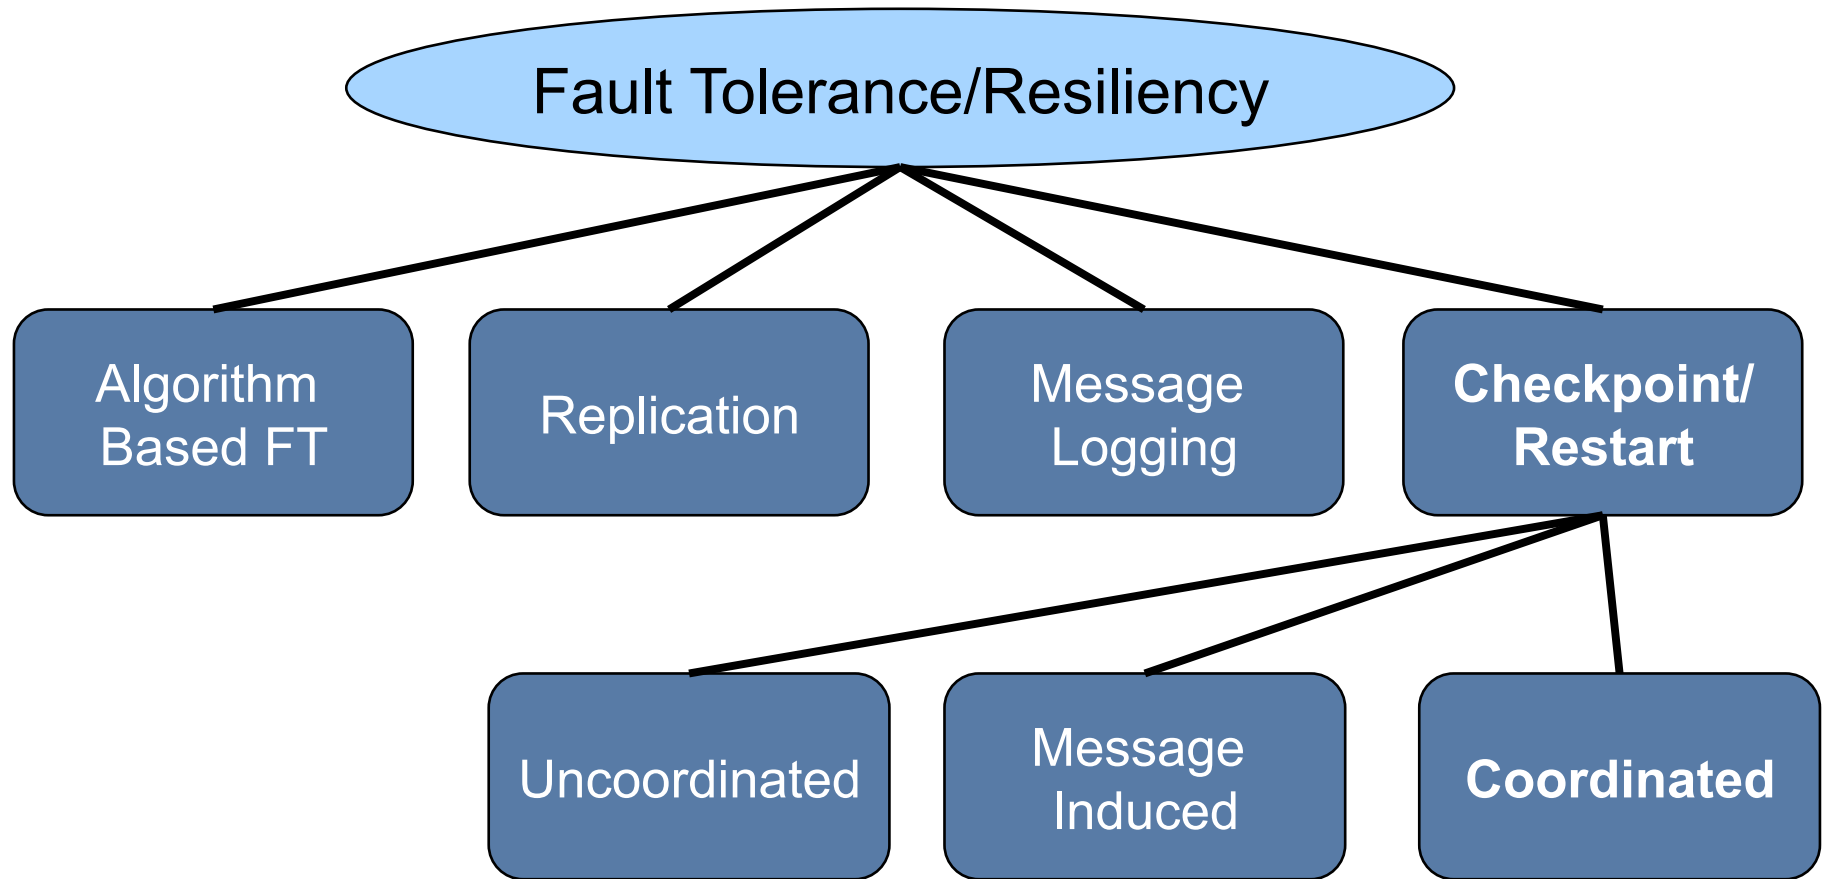


High Level Goals

- Deliver usable features to end users
 - Don't publish and run
- Extensible C/R research infrastructure
 - Focused development areas
 - Apples-to-apples comparisons
 - Opportunities for public release & support



Fault Tolerance, Debugging, Process Migration



Fault Tolerance, Debugging, Process Migration

Questions

Joshua Hursey
jjhursey@open-mpi.org

osl.iu.edu/research/ft/

www.cs.indiana.edu/~jjhursey/



OPEN MPI



INDIANA UNIVERSITY
PERVASIVE TECHNOLOGY INSTITUTE